WHAT IS CLAIMED IS:

1. For use in a switch fabric, a routing mechanism for directing data transfers through the switch fabric between an input port and an output port, wherein the switch fabric includes a plurality of paths from the input port to the output port, the routing mechanism comprising:

a controller determining, for each of the plurality of paths, a likelihood of existing traffic blocking a desired data transfer from the input port to the output port on the respective path,

wherein the controller selects one of the plurality of paths having a least likelihood of being blocked by the existing traffic for the desired data transfer.

2. The routing mechanism as set forth in Claim 1 wherein the switch fabric comprises a plurality of switches each having a plurality of input ports, at least one input queue associated with each input port, and a plurality of output ports, wherein the input and output ports of the plurality of switches are interconnected to form a multi-stage switch mesh, the routing mechanism further comprising:

a tabulation of blocking count values for each output port within the multi-stage switch mesh on one of the plurality of paths, wherein each blocking count value represents a traffic load of existing traffic through the respective output port, the controller determining the likelihood of existing traffic blocking the desired data transfer for each of the paths within the plurality of paths by summing blocking count values for all output ports on the respective path, wherein the selected one of the plurality of paths has a lowest total blocking count value.

3. The routing mechanism as set forth in Claim 2 wherein the blocking count values for each output port represent a traffic load of existing traffic through the respective output port from both the input port for each of the plurality of paths and other input ports within the switch fabric.

4. The routing mechanism as set forth in Claim 2 wherein the tabulation includes a plurality of blocking count values for each output port within the multi-stage switch mesh on one of the plurality of paths, each of the plurality of blocking count values representing a traffic load of existing traffic through the respective output port at a priority greater than or equal to one of a plurality of priorities,

wherein the controller employs a blocking count value corresponding to a priority of the desired data transfer in determining the likelihood of existing traffic blocking the desired data transfer for each of the paths within the plurality of paths.

5. The routing mechanism as set forth in Claim 4 wherein each blocking count value for each output port represents a traffic load of existing traffic through the respective output port from any input port within the switch fabric at a priority greater than or equal to a corresponding one of a plurality of priorities.

6. The routing mechanism as set forth in Claim 2 wherein the tabulation includes an existing traffic load for each of the plurality of paths,

wherein, when two or more of the plurality of paths each have a same total blocking count value which is lower than total blocking count values for all remaining paths within the plurality of paths, the controller selects one of the two or more paths having a lowest existing traffic load.

7. The routing mechanism as set forth in Claim 6 wherein controller employs the existing traffic load for each of the plurality of paths to identify a subset of paths having sufficient capacity for the desired data transfer and selects one of the subset of paths having a lowest total blocking count value as a route for the desired data transfer.

8. A switch fabric comprising:

a plurality of input ports;

a plurality of output ports, wherein each of the plurality of input ports is coupled to each of the plurality of output paths by two or more paths;

a routing mechanism for directing data transfers through the switch fabric between an input port within the plurality of input ports and an output port within the plurality of output ports, wherein the switch fabric includes a plurality of paths from the input port to the output port, the routing mechanism comprising:

a controller determining, for each of the plurality of paths, a likelihood of existing traffic blocking a desired data transfer from the input port to the output port on the respective path,

wherein the controller selects one of the plurality of paths having a least likelihood of being blocked by the existing traffic for the desired data transfer.

9. The switch fabric as set forth in Claim 8 further comprising:

a plurality of switches each having a plurality of input ports, at least one input queue associated with each input port, and a plurality of output ports, wherein the input and output ports of the plurality of switches are interconnected to form a multi-stage switch mesh, the plurality of input ports for the switch fabric formed by input ports for switches within a first stage of the multi-stage switch mesh and the plurality of output ports for the switch fabric formed by output ports for switches within a last stage of the multi-stage switch mesh;

a tabulation of blocking count values for each output port within the multi-stage switch mesh on one of the plurality of paths, wherein each blocking count value represents a traffic load of existing traffic through the respective output port, the controller determining the likelihood of existing traffic blocking the desired data transfer for each of the paths within the plurality of paths by summing blocking count values for all output ports on the respective path, wherein the selected one of the plurality of paths has a lowest total blocking count value.

10. The switch fabric as set forth in Claim 9 wherein the blocking count values for each output port represent a traffic load of existing traffic through the respective output port from both the input port for each of the plurality of paths and other input ports within the switch fabric.

11. The switch fabric as set forth in Claim 9 wherein the tabulation includes a plurality of blocking count values for each output port within the multi-stage switch mesh on one of the plurality of paths, each of the plurality of blocking count values representing a traffic load of existing traffic through the respective output port at a priority greater than or equal to one of a plurality of priorities,

wherein the controller employs a blocking count value corresponding to a priority of the desired data transfer in determining the likelihood of existing traffic blocking the desired data transfer for each of the paths within the plurality of paths.

12. The switch fabric as set forth in Claim 11 wherein each blocking count value for each output port represents a traffic load of existing traffic through the respective output port from any input port within the switch fabric at a priority greater than or equal to a corresponding one of a plurality of priorities.

13. The switch fabric as set forth in Claim 9 wherein the tabulation includes an existing traffic load for each of the plurality of paths,

wherein, when two or more of the plurality of paths each have a same total blocking count value which is lower than total blocking count values for all remaining paths within the plurality of paths, the controller selects one of the two or more paths having a lowest existing traffic load.

14. The switch fabric as set forth in Claim 13 wherein controller employs the existing traffic load for each of the plurality of paths to identify a subset of paths having sufficient capacity for the desired data transfer and selects one of the subset of paths having a lowest total blocking count value as a route for the desired data transfer.

15. For use in a switch fabric, a method of directing data transfers through the switch fabric between an input port and an output port, wherein the switch fabric includes a plurality of paths from the input port to the output port, the method comprising:

determining, for each of the plurality of paths, a likelihood of existing traffic blocking a desired data transfer from the input port to the output port on the respective path; and

selecting one of the plurality of paths having a least likelihood of being blocked by the existing traffic for the desired data transfer.

16. The method as set forth in Claim 15 wherein the switch fabric comprises a plurality of switches each having a plurality of input ports, at least one input queue associated with each input port, and a plurality of output ports, wherein the input and output ports of the plurality of switches are interconnected to form a multi-stage switch mesh, the method further comprising:

maintaining a tabulation of blocking count values for each output port within the multi-stage switch mesh on one of the plurality of paths, wherein each blocking count value represents a traffic load of existing traffic through the respective output port; and

determining the likelihood of existing traffic blocking the desired data transfer for each of the paths within the plurality of paths by summing blocking count values for all output ports on the respective path, wherein the selected one of the plurality of paths has a lowest total blocking count value.

17. The method as set forth in Claim 16 further comprising:

setting the blocking count values for each output port to represent a traffic load of existing traffic through the respective output port from both the input port for each of the plurality of paths and other input ports within the switch fabric.

18. The method as set forth in Claim 16 further comprising:

maintaining, within the tabulation, a plurality of blocking count values for each output port within the multi-stage switch mesh on one of the plurality of paths, each of the plurality of blocking count values representing a traffic load of existing traffic through the respective output port at a priority greater than or equal to one of a plurality of priorities; and

employing a blocking count value for each output port corresponding to a priority of the desired data transfer in determining the likelihood of existing traffic blocking the desired data transfer for each of the paths within the plurality of paths.

19. The method as set forth in Claim 18 further comprising:

setting each blocking count value for each output port to represent a traffic load of existing traffic through the respective output port from any input port within the switch fabric at a priority greater than or equal to a corresponding one of a plurality of priorities.

20. The method as set forth in Claim 16 further comprising:

maintaining, within the tabulation, an existing traffic load for each of the plurality of paths;

employing the existing traffic load for each of the plurality of paths to identify a subset of paths having sufficient capacity for the desired data transfer;

selecting one of the subset of paths having a lowest total blocking count value as a route for the desired data transfer; and

when two or more of the subset of paths each have a same total blocking count value which is lower than total blocking count values for all remaining paths within the subset of paths, selecting one of the two or more paths having a lowest existing traffic load.